IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

UTILITY PATENT APPLICATION

FOR

METHOD AND SYSTEM FOR PROVIDING DOCUMENT RETENTION USING CRYPTOGRAPHY

Inventor(s): Satyajit Nath

Assignee: PSS Systems, Inc.

METHOD AND SYSTEM FOR PROVIDING DOCUMENT RETENTION USING CRYPTOGRAPHY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to: (i) U.S. Patent Application No. _____, filed concurrently herewith, and entitled "METHOD AND SYSTEM FOR PROVIDING CRYPTOGRAPIC DOCUMENT RETENTION WITH OFF-LINE ACCESS," which is hereby incorporated herein by reference; (ii) U.S. Patent Application No. 10/206,737, filed July 26, 2002, and entitled "METHOD AND SYSTEM FOR UPDATING KEYS IN A DISTRIBUTED SECURITY SYSTEM," which is hereby incorporated herein by reference; (iii) U.S. Patent Application No. 10/676,850, filed September 30, 2003, and entitled "METHOD AND SYSTEM FOR SECURING DIGITAL ASSETS USING TIME-BASED SECURITY CRITERIA," which is hereby incorporated herein by reference; (iv) U.S. Patent Application No. 10/405,587, filed April 1, 2003, and entitled "METHOD AND SYSTEM FOR SECURING DIGITAL ASSETS USING CONTENT TYPE DESIGNATIONS," which is hereby incorporated herein by reference; (v) U. S. Patent Application No.: 10/246,079, filed September 17, 2002, and entitled "METHOD AND APPARATUS FOR GENERATING KEYS FROM ACCESS RULES IN A DECENTRALIZED MANNER AND METHODS THEREFOR," which is hereby incorporated herein by reference; (vi) U. S. Patent Application No.: 10/186,203, filed June 26, 2002, and entitled "METHOD AND SYSTEM FOR IMPLEMENTING CHANGES TO SECURITY POLICIES IN A DISTRIBUTED SECURITY SYSTEM," which is hereby incorporated herein by reference; (vii) U. S. Patent Application No.: 10/159,537, filed May 5, 2002, and entitled "METHOD AND APPARATUS FOR SECURING DIGITAL ASSETS," which is hereby incorporated herein by reference; and (viii) U. S. Patent Application No.: 10/127,109, filed April 22, 2002, and entitled "EVALUATION OF ACCESS RIGHTS TO SECURED DIGITAL ASSETS," which is hereby incorporated herein by reference.

BACKGROUND OF THE INVENTION

Field of the Invention

[0002] The present invention relates to security systems for data and, more particularly, to security systems that protect electronic files in an inter/intra enterprise environment.

Description of Related Art

[0003] The Internet is the fastest growing telecommunications medium in history. This growth and the easy access it affords have significantly enhanced the opportunity to use advanced information technology for both the public and private sectors. It provides unprecedented opportunities for interaction and data sharing among businesses and individuals. However, the advantages provided by the Internet come with a significantly greater element of risk to the confidentiality and integrity of information. The Internet is an open, public and international network of interconnected computers and electronic devices. Without proper security means, an unauthorized person or machine may intercept information traveling across the Internet and even gain access to proprietary information stored in computers that interconnect to the Internet.

[0004] There are many efforts in progress aimed at protecting proprietary information traveling across the Internet and controlling access to computers carrying the proprietary information. Cryptography allows people to carry over the confidence found in the physical world to the electronic world, thus allowing people to do business electronically without worries of deceit and deception. Every day millions of people interact electronically, whether it is through *e*-mail, *e*-commerce (business conducted over the Internet), ATM machines, or cellular phones. The perpetual increase of information transmitted electronically has led to an increased reliance on cryptography.

[0005] One of the ongoing efforts in protecting the proprietary information traveling across the Internet is to use one or more cryptographic techniques to secure a private communication session between two communicating computers on the Internet. The cryptographic techniques provide a way to transmit information across an unsecure communication channel without disclosing the contents of the information to anyone eavesdropping on the communication channel. Using an

encryption process in a cryptographic technique, one party can protect the contents of the data in transit from access by an unauthorized third party, yet the intended party can read the encrypted data after using a corresponding decryption process.

[0006] A firewall is another security measure that protects the resources of a private network from users of other networks. However, it has been reported that many unauthorized accesses to proprietary information occur from the inside, as opposed to from the outside. An example of someone gaining unauthorized access from the inside is when restricted or proprietary information is accessed by someone within an organization who is not supposed to do so. Due to the open nature of networks, contractual information, customer data, executive communications, product specifications, and a host of other confidential and proprietary intellectual property remain available and vulnerable to improper access and usage by unauthorized users within or outside a supposedly protected perimeter.

[0007] Many businesses and organizations have been looking for effective ways to protect their proprietary information. Typically, businesses and organizations have deployed firewalls, Virtual Private Networks (VPNs) and Intrusion Detection Systems (IDS) to provide protection. Unfortunately, these various security means have been proven insufficient to reliably protect proprietary information residing on private networks. For example, depending on passwords to access sensitive documents from within often causes security breaches when the password of a few characters long is leaked or detected. Consequently, various cryptographic means are deployed to provide restricted access to electronic data (namely, electronic documents) in security systems.

[0008] Besides restricting access to electronic documents, businesses and organizations also face the difficulty of implementing document retention for electronic documents. In today's heavily regulated business environment, various electronic documents need to be retained for a certain period of time and thereafter may be disposed of. Although conventional approaches are able to retain documents for a period of time and then dispose of them, these conventional approaches require that the retention duration be known and specified. With file security systems that oversee the creation and securing of electronic documents, often the electronic documents have retention periods that depend on future events that are unscheduled. Unfortunately, however, conventional approaches are unable

to adequately handle document retention when unscheduled future events are involved.

[0009] Therefore, there is a need to provide more effective ways to impose document retention restrictions on electronic resources even when unscheduled future events are involved.

SUMMARY OF THE INVENTION

[0010] The invention relates to techniques for utilizing security criteria to implement document retention for electronic documents. The security criteria can also limit when, how and where access to the electronic documents is permitted. The security criteria can pertain to keys (or ciphers) used to secure (e.g., encrypt) electronic files (namely, electronic documents), or to unsecure (e.g., decrypt) electronic files already secured. At least a portion of the security criteria can be used to implement document retention, namely, a document retention policy. After a secured electronic document has been retained for the duration of the document retention policy, the associated security criteria becomes no longer available, thus preventing subsequent access to the secured electronic document. In other words, access restrictions on electronic documents can be used to prevent access to electronic documents which are no longer to be retained.

[0011] In one embodiment, the security criteria can be managed by a document retention system. In another embodiment, the security criteria can be managed more generally by a file security system.

[0012] The invention can be implemented in numerous ways, including as a method, system, device, and computer readable medium. Several embodiments of the invention are discussed below.

[0013] As a method of providing automated document retention for electronic documents, one embodiment of the invention includes the acts of: obtaining an electronic document; assigning a document retention policy to the electronic document, the document retention policy being based on a future event that is unscheduled; and cryptographically imposing the document retention policy on the electronic document.

[0014] As a method for restricting access to an electronic document, one embodiment of the invention includes the acts of: identifying an electronic document to be secured, the electronic document having at least a data portion that contains data; obtaining a document key; encrypting the data portion of the electronic document using the document key to produce an encrypted data portion; obtaining a retention access key, the retention access key being used to enforce a document retention policy on the electronic document; encrypting the document key using the retention access key to produce an encrypted document key; forming a secured electronic document from at least the encrypted data portion and the encrypted document key; and storing the secured electronic document.

[0015] As a method for accessing a secured electronic document by a requestor, the secured electronic document having at least a header portion and a data portion, one embodiment of the invention includes the acts of: obtaining a retention access key, the retention access key being used to enforce a document retention policy on the electronic document; obtaining an encrypted document key from the header portion of the secured electronic document; decrypting the encrypted document key using the retention access key to produce a document key; decrypting an encrypted data portion of the secured electronic document using the document key to produce a data portion; and supplying the data portion to the requestor.

[0016] As a method for distributing cryptographic keys used in a file security system, one embodiment of the invention includes the acts of: receiving a request for a document retention key that is necessary to gain access to a cryptographically secured electronic document; identifying a document retention period associated with the document retention key, the document retention period being dependent on a future event that was unscheduled when the document retention period was associated with the electronic document; determining whether the document retention period associated with the document retention key has been exceeded; and refusing to distribute the document retention key in response to the request when it is determined that the document retention period for the electronic document has been exceeded.

[0017] As a file security system for restricting access to electronic files, one embodiment of the invention includes at least a key store and an access manager operatively connected to the key store. The key store stores a plurality of

cryptographic key pairs. Each of the cryptographic key pairs includes a public key and a private key, and at least one of the cryptographic key pairs pertains to a retention policy that is dependent on a future event. The access manager determines whether the private key of the at least one of the cryptographic key pairs pertaining to the retention policy is permitted to be provided to a requestor based on whether the future event has occurred. The requestor requires the private key of the at least one of the cryptographic key pairs pertaining to the retention policy to access a secured electronic file. The secured electronic file was previously secured using the public key of the at least one of the cryptographic key pairs pertaining to the retention policy, and the future event was unscheduled at the time the electronic file was so secured.

[0018] As a computer readable medium including at least computer program code for providing automated data retention for electronic data, one embodiment of the invention includes at least: computer program code for obtaining electronic data; computer program code for assigning a data retention policy to the electronic data, the data retention policy being based on a future event that is unscheduled; and computer program code for cryptographically imposing the data retention policy to the electronic data.

[0019] Other objects, features, and advantages of the present invention will become apparent upon examining the following detailed description of an embodiment thereof, taken in conjunction with the attached drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

[0021] FIG. 1 is a block diagram of a file security system according to one embodiment of the invention.

[0022] FIG. 2 is a flow diagram of a retention policy assignment process according to one embodiment of the invention.

[0023] FIG. 3 is a flow diagram of a future event evaluation process according to one embodiment of the invention.

[0024] FIG. 4 is a flow diagram of an expiration process according to one embodiment of the invention.

[0025] FIG. 5 is a flow diagram of an access request process according to one embodiment of the invention.

[0026] FIG. 6 is a flow diagram of a file securing process according to one embodiment of the invention.

[0027] FIGs. 7A and 7B are flow diagrams of a document securing process according to one embodiment of the invention.

[0028] FIG. 8 is a flow diagram of a document unsecuring process according to one embodiment of the invention.

[0029] FIG. 9 is a flow diagram of an access key retrieval process according to one embodiment of the invention.

[0030] FIG. 10 shows a basic security system in which the invention may be practiced in accordance with one embodiment thereof.

[0031] FIG. 11 shows an exemplary data structure of a secured file that may be used in one embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

[0032] The invention relates to techniques for utilizing security criteria to implement document retention for electronic documents. The security criteria can also limit when, how and where access to the electronic documents is permitted. The security criteria can pertain to keys (or ciphers) used to secure (e.g., encrypt) electronic files (namely, electronic documents), or to unsecure (e.g., decrypt) electronic files already secured. At least a portion of the security criteria can be used to implement document retention, namely, a document retention policy. After a secured electronic document has been retained for the duration of the document retention policy, the associated security criteria becomes no longer available, thus preventing subsequent access to the secured electronic document. In other words,

access restrictions on electronic documents can be used to prevent access to electronic documents which are no longer to be retained.

[0033] In one embodiment, the security criteria can be managed by a document retention system. In another embodiment, the security criteria can be managed more generally by a file security system.

[0034] Secured files are files that require one or more keys, passwords, access privileges, etc. to gain access to their content. The security is often provided through encryption and access rules. The files, for example, can pertain to documents, multimedia files, data, executable code, images and text. In general, a secured file can only be accessed by authenticated users with appropriate access rights or privileges. In one embodiment, each secured file is provided with a header portion and a data portion, where the header portion contains, or points to, security information. The security information is used to determine whether access to associated data portions of secured files is permitted.

[0035] As used herein, a user may mean a human user, a software agent, a group of users, a member of the group, a device and/or application. Besides a human user who needs to access a secured document, a software application or agent sometimes needs to access secured files in order to proceed. Accordingly, unless specifically stated, the "user" as used herein does not necessarily pertain to a human being.

[0036] The invention is related to processes, systems, architectures and software products for providing automated retention of digital assets (e.g., electronic documents). The invention is particularly suitable in an enterprise environment. The invention can also be implemented by a security system that additionally secures digital assets (i.e., secured data) so that only authenticated users with appropriate access rights or privileges can gain access thereto. Digital assets may include, but not be limited to, various types of documents, multimedia files, data, executable code, images and text.

[0037] In the following description, numerous specific details are set forth in order to provide a thorough understanding of the invention. However, it will become obvious to those skilled in the art that the invention may be practiced without these specific details. The description and representation herein are the common

meanings used by those experienced or skilled in the art to most effectively convey the substance of their work to others skilled in the art. In other instances, well-known methods, procedures, components, and circuitry have not been described in detail to avoid unnecessarily obscuring aspects of the invention.

[0038] Reference herein to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in connection with the embodiment can be included in at least one embodiment of the invention. The appearances of the phrase "in one embodiment" in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Further, the order of blocks in process flowcharts or diagrams representing one or more embodiments of the invention do not inherently indicate any particular order, nor imply any limitations in the invention.

[0039] Embodiments of the invention are discussed herein with reference to FIGs. 1 – 11. However, those skilled in the art will readily appreciate that the detailed description given herein with respect to these figures is for explanatory purposes as the invention extends beyond these limited embodiments.

[0040] FIG. 1 is a block diagram of a file security system 100 according to one embodiment of the invention. The file security system 100 includes an access server 102 that provides central management for the file security system 100. The access server 102 can store or retrieve files from a server file store 104. The access server 102 can control the distribution of cryptographic keys from a key store 106. In addition, the access server 102 can generate cryptographic keys that are stored to the key store 106. Alternatively, the access server 102 can store cryptographic keys in the key store 106 that have been received by the access server 102.

[0041] The file security system 100 also includes user machines 108 and user file stores 112. The user machines 108 couple to the access server 102 via a network 110. The network 110 can be a private network or a public network. The user machine 108 also has a user file store 112 coupled thereto. The user file store 112 can store electronic files locally for the user of the corresponding user machine 108. On the other hand, the server file store 104 can provide centralized, remote storage of electronic files for any of the users of the user machines 108.

[0042] According to the invention, the file security system 100 enables a user at the user machine 108 to secure an electronic file (document) such that access to the secured electronic file is restricted. In one embodiment of the invention, the access restriction is or includes a document retention restriction. As an example, the document retention restriction could permit subsequent access to the secured electronic file only during the retention period. After the retention period, the secured electronic file would be unable to be unsecured and thus subsequent access would be cryptographically prevented. In one embodiment, the retention period is initially based on a future event that is unscheduled. Here, the file security system 100 can interact with an event evaluator 114 to determine (e.g., periodically) whether the future event has now been scheduled. Once the future event is scheduled, the retention period is determinable. The event evaluator 114 can be referred to as a remote resource that is used to evaluate future events. The event evaluator 114 can be a web server or a document management system (e.g., a contract management system).

[0043] Once an electronic file has been secured, a user at a user machine can attempt to access the secured electronic file. In doing so, the user machine for the user would need to access the access server 102 and retrieve the appropriate one or more cryptographic keys from the key store 106 that are needed to unsecure the secured electronic file. However, after expiration of the retention period for the secured electronic document, the access server 102 does not permit the delivery of at least certain cryptographic keys from the key store 106 (private keys) that are needed to unsecure secured electronic files. The access server 102 normally also requires the user to be authorized to access the electronic file prior to delivery of at least certain cryptographic keys from the key store 106. Typically, the cryptographic keys needed to unsecure a secured electronic file are private keys.

[0044] Nevertheless, once the restrictions and/or authorizations have been satisfied and the private keys have been supplied, the private keys are usable to unsecure the secured electronic files. Typically, the private keys will expire after a relatively short period of time (e.g., 1-30 days) so that users are forced to again retrieve the necessary cryptographic keys.

[0045] FIG. 2 is a flow diagram of a retention policy assignment process 200 according to one embodiment of the invention. The retention policy assignment

process 200 operates to cryptographically secure an electronic document to implement document retention. The retention policy assignment process 200 is, for example, performed by a computing device, such as the access server 102 or the user machines 108 illustrated in FIG. 1.

[0046] The retention policy assignment process 200 initially opens or creates 200 an electronic document. A user of a computing device may assist with the opening or creation of the electronic document. Next, a decision 202 determines whether document retention is requested. Here, according to the invention, document retention policies can be imposed on the electronic document. Hence, the decision 202 determines whether document retention is to be imposed on the electronic document. When the decision 202 determines that document retention is not to be imposed, then the retention policy assignment process 200 is complete and ends with no document retention policy being imposed.

[0047] On the other hand, when the decision 202 determines that document retention is requested, then a document retention policy is specified 204 based on a future event that is presently unscheduled. Typically, the document retention policy specifies that the electronic document is to be maintained until some future point in time. However, when the document retention policy is based on a future event which is presently unscheduled, the future point in time is not known and cannot be determined. Once the future event becomes scheduled, then the period of time for document retention can be determined. In other words, the document retention policy becomes determinable because the future event is no longer unscheduled. Next, the document retention policy is cryptographically imposed 206 on the electronic document. Recall, however, that the document retention policy at this point is based on a future event which is presently unscheduled. In one implementation, a cryptographic key is utilized to encrypt the electronic document so that access to the electronic document can be restricted after the document retention policy has been exceeded. In other words, after the period of time for document retention specified by the document retention policy has been exceeded, the cryptographic key that is needed to gain access to the electronic document is no longer made available to users. As a result, because the electronic document was previously cryptographically secured using a cryptographic key, if the corresponding or counterpart cryptographic key is no longer available, then the electronic document

remains encrypted and thus unusable. In any case, following the operation 206, the retention policy assignment process 200 is complete and ends.

[0048] Once a document retention policy has been assigned to an electronic document, the document retention system or file security system 100 as shown in FIG. 1 needs to periodically evaluate whether future events associated with the document retention policies are now scheduled. In one embodiment, the access server 102 shown in FIG.1 can be utilized to periodically evaluate whether future events have become scheduled.

[0049] FIG. 3 is a flow diagram of a future event evaluation process 300 according to one embodiment of the invention. The future event evaluation process 300 can, for example, be performed by the access server 102 illustrated in FIG. 1.

[0050] The future event evaluation process 300 initially identifies 302 a future event. Here, the future event is a future event that is being monitored by a document retention system (or file security system) because it is utilized by one or more document retention policies assigned to one or more electronic documents managed by the system. After the future event has been identified 302, a remote resource can be queried 304 for status of the future event. The remote resource can, for example, be a file, a web server or an external system. An example of an external system would be a document management system or a contract management system. In any case, the event evaluator 114 shown in FIG. 1 can represent the remote resource that is being queried 304.

[0051] Next, a decision 306 determines whether a status response has been received from the remote resource. When the decision 306 determines that a status response has not been received, then the future event evaluation process 300 awaits a status response. However, after a period of time in which no response is received, the status response can be deemed or default to indicate that the future event remains unscheduled. On the other hand, when the decision 306 determines that the status response has been received, a decision 308 determines whether the future event is now scheduled by examination of the status response. When the decision 308 determines that the future event is now scheduled, then schedule information pertaining to the future event can be stored 310. For example, the access server 102 shown in FIG. 1 can store schedule information for the future

event. At a minimum, the schedule information can contain an indication that the future event is now scheduled. Typically, the schedule information would specify a date representing the occurrence (past or future) of the future event. Alternatively, when the decision 308 determines that the future event still remains unscheduled, then the operation 310 is bypassed.

[0052] Following the operation 310 or its being bypassed, a decision 312 determines whether there are other future events to be evaluated. When the decision 312 determines that there are other future events to be evaluated, the future event evaluation process 300 returns to repeat the operation 302 and subsequent operations, thereby allowing other future events to be similarly evaluated. On the other hand, when the decision 312 determines that there are no other future events to be evaluated, then the future event evaluation process 300 is complete and ends. Typically, the future event evaluation process 300 would be periodically invoked to evaluate whether any future events have become scheduled.

[0053] In one embodiment, when the remote resource can be addressed for requests (queries) by a Universal Resource Locator (URL). The URL could point to a file, a web-server or some other web-based application. In case where the URL points to a file, the file stores and can provide the status response as to whether the associated future event is unscheduled or not. The URL specifies the correct file and can do so by identifying the descriptions of future events, an event type and an event identifier.

[0054] In another embodiment, when the remote resource is a web server or an external system, the request (query) to the web server or external system can use a URL to access the web server. The URL can specify the web server or external system and describe the future event of interest. Alternatively, the query or request to the web server or external system can be a markup language (e.g., XML) document. Such a status request would also at least describe the future event of interest. The status response from the web server or external system can return an indication as to whether the associated future event is unscheduled or not. As an example, the status response can be a markup language (e.g., XML) document. It should be noted that the status response can also be signed with an electronic signature that can be used to authenticate its origin.

[0055] In still another embodiment, the remote resource is network accessible (e.g., web server or external system). Here, the remote resource can be accessed using networking techniques, such as TCP/IP networks, to get future event information from the remote resource.

[0056] FIG. 4 is a flow diagram of an expiration process 400 according to one embodiment of the invention. The expiration process 400 represents other processing that determines whether document retention periods associated with electronic documents have been exceeded, and if so, renders the associated electronic documents inaccessible. The expiration process 400 would typically be periodically invoked. The expiration process 400 can, for example, be performed by the access server 102 illustrated in FIG. 1.

[0057] The expiration process 400 initially identifies 402 a future event. Here, the future event is a future event associated with a document retention policy that is being utilized to retain one or more electronic documents by a document retention system (or file security system).

[0058] Next, a decision 404 determines whether the future event has been scheduled. As an example, the decision 404 can examine schedule information that is stored by the future event evaluation process 300 illustrated in FIG. 3. In any case, the decision 404 determines whether the future event has now been scheduled. When the decision 404 determines that the future event has been scheduled, then a decision 406 determines whether a document retention period associated with the now scheduled future event has expired. For example, a document retention period might typically be represented as a predetermined period of time following a future event. Hence, once the future event is scheduled and thus has a date certain, the document retention period is determinable.

[0059] When the decision 406 determines that the document retention period has expired, then the cryptographic key associated with the document retention policy is identified 408. The document retention policy may be associated with one or a plurality of cryptographic keys that are utilized to secure one or a plurality of different electronic documents. In any event, once the cryptographic key is identified 408, the cryptographic key is deactivated 410. In other words, the cryptographic key utilized to implement the document retention policy for the electronic document is destroyed,

deleted or disabled. Consequently, the cryptographic key is no longer useable to gain access to the electronic document that has been encrypted therewith, thereby implementing the document retention policy. In other words, the associated electronic document is thereafter inaccessible by those persons or machines that were previously able to access the electronic document. In effect, the electronic document has been effectively destroyed. The operations 408 and 410 are bypassed when the decision 404 determines that the future event has not yet been scheduled or when the decision 406 determines that the document retention period has not yet expired.

[0060] In any case, following the operation 410 or its being bypassed, a decision 412 determines whether there are other future events to be similarly processed so as to determine whether the associated document retention period or periods have expired. Hence, when the decision 412 determines that other future events are to be processed, the expiration process 400 returns to repeat the operation 402 and subsequent operations to process another future event.

[0061] FIG. 5 is a flow diagram of an access request process 500 according to one embodiment of the invention. The access request process 500 is, for example, performed by a computing device, such as the user machines 108 illustrated in FIG. 1.

[0062] The access request process 500 begins with a decision 502 that determines whether an electronic document access request has been received. A user of a computing device can initiate an electronic document access request. When the decision 502 determines that an electronic document access request has not been received, then the access request process 500 awaits such as request. Once the decision 502 determines that a document access request has been received, then a decision 504 can determine whether document retention is imposed on the electronic document that is to be accessed. When the decision 504 determines that document retention is not imposed on the electronic document to be accessed, access to the electronic document is permitted 506.

[0063] On the other hand, when the decision 504 determines that document retention is imposed on the electronic document to be accessed, a cryptographic key associated with the document retention policy that is imposed on the electronic

document is requested 508. A decision 510 then determines whether the requested key has been received. Here, the requested key is the cryptographic key that has been requested 508. In one implementation, such as shown in FIG. 1, the user machine 108 requests the key from the access server 102, and the key is provided (if at all) to the user machine 108 via the network 110.

[0064] When the decision 510 determines that the requested key has not been received, then access to the electronic document is denied 512. In this case, the document retention policy causes the cryptographic key to be no longer available to the requestor. In such case, although the requestor may have access to the electronic document, since the cryptographic key is not available, the requestor is not able to gain access to the electronic document. In other words, the electronic document remains in its encrypted format and thus unusable by the requestor. In such case, the document retention policy imposed on the electronic document caused the electronic document to expire.

[0065] On the other hand, when the decision 510 determines that the requested key has been received, then access to the electronic document is permitted 514 through use of the cryptographic key. In other words, the cryptographic key can be used to decrypt the encrypted electronic document, thereby allowing the requestor to gain access to the electronic document.

[0066] Following the operations 506, 512 and 514, the access request process 500 is complete and ends. However, it should be noted that additional layers of encryption could be imposed on the electronic document besides the level of encryption utilized to implement a document retention policy. Hence, other keys or requirements can be imposed by a file security system in order to further restrict access to the electronic documents. For example, co-pending U.S. Patent Application No. 10/405,587, filed April 1, 2003 and entitled "METHOD AND SYSTEM FOR SECURING DIGITAL ASSETS USING CONTENT TYPE DESIGNATIONS," describes representative other document security measures and is hereby incorporated herein by reference.

[0067] As previously noted, document retention can be provided by a file security system. In which case, an electronic file can be cryptographically secured using a key that is associated with file retention. The key can then automatically be made

unavailable after the retention period for the electronic document has been exceeded. Accordingly, the electronic file is no longer accessible after the retention period is exceeded. Typically, such inaccessibility is permanent and irreversible. However, in special situations, the automatic nature of the non-retention can be suspended so as to permit extended access. For example, a system administrator can cause the retention period to be extended so that the electronic file remains accessible even after the original extension period has been scheduled. This override capability can be important when, for example, legal issues arise concerning the electronic file. In such cases, the electronic file should be retained and not automatically expired. In one embodiment, the system administrator can either override an assigned retention period or set a new retention period (which may also be based on a future event).

[0068] FIG. 6 is a flow diagram of a file securing process 600 according to one embodiment of the invention. The file securing process 600 is, for example, performed by a computing device, such as the access server 102 or the user machines 108 illustrated in FIG. 1.

[0069] The file securing process 600 initially identifies 602 an electronic file to be secured. Here, the electronic file is identified to the computing device carrying out the file securing process 600. A user of the computing device may assist in the identification of the electronic file. After the electronic file to be secured has been identified 602, a document retention access key is obtained 604. Typically, the document retention access key is obtained 604 from an access server. For example, if the file securing process 600 is performed by the user machine 108, the document retention access key can be retrieved remotely from the access server 102, which can retrieve the document retention access key from the key store 106. Here, the document retention access key is a public key and is normally readily available.

[0070] Next, after the time-based access key has been obtained 604, the electronic file is secured 606 through use of the document retention access key. The result of the securing of the electronic file is to produce a secured electronic file. Typically, the electronic file is cryptographically secured through encryption (directly or indirectly) using the document retention access key. In one embodiment, one or more other keys are used to further control access to the electronic file. Thereafter,

the secured electronic file is stored 608. After the secured electronic file is stored 608, the file securing process 600 ends. Following the securing process 600, the secured electronic file can be distributed to one or more user machines 108 having interested users.

[0071] FIGs. 7A and 7B are flow diagrams of a document securing process 700 according to one embodiment of the invention. The document securing process 700 is, for example, performed by a computing device, such as the access server 102 or the user machines 108 illustrated in FIG. 1.

[0072] The document securing process 700 opens or creates 702 an electronic document. At this point, the electronic document is unsecure, which is often referred to as being in the "clear." Next, a decision 704 determines whether the electronic document is to be secured. Here, the user or creator of the electronic document has the option of securing the document, if so desired. When the decision 704 determines that the electronic document is not to be secured, then the electronic document is saved 706. Here, the electronic document being saved is not secured. Following the block 706, the document securing process 700 ends with the electronic document being saved in an unsecured fashion.

[0073] On the other hand, when the decision 704 determines that the electronic document is to be secured, then a data portion of the electronic document is encrypted 708 using a document key. The document key can be a cryptographic key that is generated or obtained. Typically, each document managed and secured by the file (document) security system would be encrypted 708 using a different document key. After the data portion of the electronic document has been encrypted 708, a decision 710 determines whether a document retention restriction should be imposed on the electronic document. The user or creator of the electronic document can have the option of securing the document with a document retention access restriction, if so desired.

[0074] When the decision 710 determines that a document retention restriction should be imposed on the electronic document, a public document retention access key is requested 712. In one embodiment, the public document retention access key can be requested from the access server 102 by the user machine 108. The access server 102 can then retrieve or generate the public document retention access key

and supply it to the user machine 108. In an alternative implementation, the user machine may have already received the public document retention access key (e.g., such as in a key cache) and thus would not need to request such.

[0075] Next, a decision 714 determines whether the public document retention access key has been received. Once the decision 714 determines that the public document retention access key has been received (or already available), the document key is encrypted 716 using the public document retention access key. Here, the document key is being encrypted using the public document retention access key. In other words, the public document retention access key is indirectly used to encrypt the electronic document by encryption of the document key. Next, a secured electronic document is formed 718 from the encrypted data portion and the encrypted document key. Thereafter, the secured electronic document is saved 720. In this case, following the block 720, the document securing process 700 ends with the electronic document being saved in a secured fashion with at least a document retention access restriction.

[0076] Alternatively, when the decision 710 determines that a document retention access restriction is not to be imposed on the electronic document, then the blocks 712-716 are bypassed. In such case, the secured electronic document is formed 718 from the encrypted data portion and the document key. Here, the document key is not encrypted using a public document retention access key. The resulting secured electronic document is then saved 720. In this case, following the block 720, the document securing process 700 ends with the electronic document being saved in a secured fashion without any time-based access restrictions.

[0077] FIG. 8 is a flow diagram of a document unsecuring process 800 according to one embodiment of the invention. The document unsecuring process 800 can be performed at a client machine or a server machine, such as the user machine 108 or the access server 102 illustrated in FIG. 1.

[0078] The document unsecuring process 800 begins with a decision 802 that determines whether a request to access a secured electronic document has been received. When the decision 802 determines that a request to access a secured electronic document has not yet been received, the document unsecuring process 800 awaits such a request. In other words, the document unsecuring process 800

can be considered to be invoked once access to a secured electronic document is requested.

[0079] Once the decision 802 determines that a request to access a secured electronic document has been received, a decision 804 determines whether a document retention restriction is present. In one implementation, the decision 804 can evaluate a header portion of the secured electronic document to determine whether a document retention restriction is present. In another implementation, the decision 804 can evaluate a system policy to determine whether a document retention restriction is present. As an example, the header can include an indicator of a document retention restriction.

[0080] When the decision 804 determines that a document retention restriction is present, then a private document retention access key is requested 806. In one embodiment, the private document retention access key is requested 806 from a file security system, such as a server machine thereof (e.g., access server 102). Then, a decision 808 determines whether the requested key has been received. When the decision 808 determines that the requested key has not yet been received, a decision 810 determines whether access to the requested key has been denied. Typically, the private document retention access key is only available so long as a retention period for the secured electronic document has not been exceeded. In one embodiment, the access server 102 controls access to the private document retention access key which is stored in the key store 106. Hence, in such an embodiment, the access server 102 would deny any request for the document retention access key after the retention period has been exceeded. In any case, when the decision 810 determines that access to the requested key has been denied, then access to the secured electronic document is denied and notice that access has been denied is returned 812. In one embodiment, the notice can more specifically indicate that access is denied because the document has expired. Following the block 812, the document unsecuring process 800 ends with access to the secured electronic document being denied.

[0081] On the other hand, when the decision 810 determines that access to the requested key has not been denied, then the document unsecuring process 800 returns to repeat the decision 808 so as to wait for the requested key to be received. Once the decision 808 determines that the requested key (the private document

retention access key) has been received, the encrypted document key from the secured electronic document is decrypted 814 using the private document retention access key to yield the document key (unencrypted). Here, in one embodiment, a header portion of the secured electronic document includes at least the encrypted document key (as well as the indicator for the private document retention access key). Next, an encrypted data portion of the secured electronic document is decrypted 816 using the document key. Finally, the data portion of the electronic document is then returned 818 to the requestor. Additionally, it should be noted that when the decision 804 determines that a document retention access restriction is not present, then the document unsecuring process 800 skips blocks 806-814 and proceeds to block 816. Following block 818, the document unsecuring process 800 ends with access to the secured electronic document being successful.

[0082] In one embodiment, the document retention access keys (e.g., the public and private key pair) can be unique (i.e., different) for each electronic document. Alternatively, to manage the number of key pairs, the document retention access keys can be shared by electronic documents being retained for a like duration.

[0083] FIG. 9 is a flow diagram of an access key retrieval process 900 according to one embodiment of the invention. The access key retrieval process 900 is, for example, performed by a server machine, such as the access server 102 illustrated in FIG. 1.

[0084] The access key retrieval process 900 begins with a decision 902 that determines whether a request for a document retention access key has been received. When the decision 902 determines that a request for a document retention access key has not yet been received, the access key retrieval process 900 awaits such a request. Once the decision 902 determines that a document retention access key has been received, the access key retrieval process 900 continues. In other words, the access key retrieval process 900 can be deemed invoked when a request for a document retention access key is received.

[0085] In any case, once the access key retrieval process 900 continues, a decision 904 determines whether the requested access key is a private key. When the decision 904 determines that the requested key is not a private key (i.e., is a public key), then a public document retention access key (which was requested) is

sent 906. Typically, the public document retention access key would be sent to a requestor (such as a user machine). In one embodiment, the public document retention access key is retrieved from a remote key store by a server and sent by the server to the requestor.

[0086] On the other hand, when the decision 904 determines that the requested key is a private key (i.e., a private document retention access key), a decision 912 determines whether the private document retention access key is available. When the decision 912 determines that the private document retention access key is not available, then the key request is denied 914. In such case, the requestor would not be able to utilize the electronic documents that have been secured with the associated document retention policy. In effect, the electronic documents would be deemed expired (i.e., no longer retained).

[0087] Alternatively, when the decision 912 determines that the private document retention access key is available, then the private document retention access key is sent 916 to the requestor. In this case, the requestor is able to use the requested key to gain access to the electronic documents. Following the blocks 906, 914 and 916, the access key retrieval process 900 ends.

[0088] FIG. 10 shows a basic security system 1000 in which the invention may be practiced in accordance with one embodiment thereof. The security system 1000 may be employed in an enterprise or inter-enterprise environment. It includes a first server 1006 (also referred to as a central server) providing centralized access management for the enterprise. The first server 1006 can control restrictive access to files secured by the security system 1000 as well as file (e.g., document) retention. To provide dependability, reliability and scalability of the system, one or more second servers 1004 (also referred to as local servers, of which one is shown) may be employed to provide backup or distributed access management for users or client machines serviced locally. The server 1004 is coupled to a network 1008 and a network 1010. For illustration purposes, there are two client machines 1001 and 1002 being serviced by the local server 1004. Alternatively, one of the client machines 1001 and 1002 may be considered as a networked storage device.

[0089] Secured files may be stored in any one of the devices 1001, 1002, 1004 and 1006. When a user of the client machine 1001 attempts to exchange a secured

file with a remote destination 1012 being used by an external user, the processes discussed above can be utilized to ensure that the requested secure file is delivered without compromising the security imposed on the secured file.

[0090] According to one embodiment, a created document is caused to go through an encryption process that is preferably transparent to a user. In other words, the created document is encrypted or decrypted under the authoring application so that the user is not aware of the process. One or more keys, such as a user key and a document retention access key, can be used to retrieve a file key to decrypt an encrypted document. Typically, the user key is associated with an access privilege for the user or a group of users, and the document retention access key is associated with a retention period imposed on the created document. For a given secured document, only a user with proper access privileges can access the secured document and then only after a time restriction, if present, is satisfied.

[0091] In one setting, a secured document may be uploaded via the network 1010 from the client computer 1001 to a computing or storage device 1002 that may serve as a central repository. Although not necessary, the network 1010 can provide a private link between the computer 1001 and the computing or storage device 1002. Such link may be provided by an internal network in an enterprise or a secured communication protocol (e.g., VPN and HTTPS) over a public network (e.g., the Internet). Alternatively, such link may simply be provided by a TCP/IP link. As such, secured documents on the computing or storage device 1002 may be remotely accessed.

[0092] In another setting, the computer 1001 and the computing or storage device 1002 are inseparable, in which case the computing or storage device 1002 may be a local store to retain secured documents or receive secured network resources (e.g., dynamic Web contents, results of a database query, or a live multimedia feed). Regardless of where the secured documents or secured resources are actually located, a user, with proper access privileges and within retention periods, can access the secured documents or resources from the client computer 1001 or the computing or storage device 1002 using an application (e.g., Microsoft Internet Explorer, Microsoft Word or Adobe Acrobat Reader).

[0093] Accordingly, respective local modules in local servers, in coordination with the central server, form a distributed mechanism to provide not only distributed access control enforcement but also file (e.g., document) retention. Such distributed access control enforcement ensures the dependability, reliability and scalability of centralized access control management undertaken by the central server for an entire enterprise or a business location.

[0094] FIG. 11 shows an exemplary data structure 1120 of a secured file that may be used in one embodiment of the invention. The data structure 1120 includes two portions: a header (or header portion) 1122 and encrypted data (or an encrypted data portion) 1124. The header 1122 can be generated in accordance with a security template associated with a data store and thus provides restrictive access to the data portion 1124 which, for example, is an encrypted version of a plain file. Optionally, the data structure 1120 may also include an error-checking portion 1125 that stores one or more error-checking codes, for example, a separate error-checking code for each block of encrypted data 1124. These error-checking codes may also be associated with a Cyclical Redundancy Check (CRC) for the header 1122 and/or the encrypted data 1124. The header 1122 includes a flag bit or signature 1127 and security information 1126 that is in accordance with the security template for the data store. According to one embodiment, the security information 1126 is encrypted and can be decrypted with a user key associated with an authenticated user (or requestor).

[0095] The security information 1126 can vary depending upon implementation. However, as shown in FIG. 11, the security information 1126 includes a user identifier (ID) 1128, access policy (access rules) 1129, keys (cryptographic keys) 1130, and other information 1131. Although multiple user identifiers may be used, a user identifier 1128 is used to identify a user or a group that is permitted to access the secured file. The access rules 1129 provide restrictive access to the encrypted data portion 1124. The keys 1130 are cipher keys (and/or pointers or indentifiers therefor) that, once obtained, can be used to decrypt the encrypted data portion 1124 and thus, in general, are protected. In one implementation of the data structure 1120, at least one of the keys 1130 is encrypted in conjunction with the access rules 1129. In another implementation of the data structure 1120, at least one of the keys 1130 is a file retention access key or is a key encrypted with a file retention access

key, either of which can possibly be further protected by the access rules 1129. The other information 1131 is an additional space for other information to be stored within the security information 1126. For example, the other information 1131 may be used to include other information facilitating secure access to the secured file, such as version number or author identifier.

[0096] The invention is preferably implemented by software or a combination of hardware and software, but can also be implemented in hardware. The invention can also be embodied as computer readable code on a computer readable medium. The computer readable medium is any data storage device that can store data which can thereafter be read by a computer system. Examples of the computer readable medium include read-only memory, random-access memory, CD-ROMs, DVDs, magnetic tape, optical data storage devices, and carrier waves. The computer readable medium can also be distributed over network-coupled computer systems so that the computer readable code is stored and executed in a distributed fashion.

[0097] The various embodiments, implementations and features of the invention noted above can be combined in various ways or used separately. Those skilled in the art will understand from the description that the invention can be equally applied to or used in various other settings with respect to different combinations, embodiments, implementations or features as provided in the description herein.

[0098] The advantages of the invention are numerous. Different embodiments or implementations may yield one or more of the following advantages. One advantage of the invention is that retention of electronic data (e.g., electronic documents) is provided in a robust, cryptographically secure manner. Another advantage of the invention is that retention policies can be based on future events that are unscheduled when assigned to electronic data. Still another advantage of the invention is that the needed cryptographic keys to unsecure secured electronic data are no longer released by a server to a client once a retention policy has expired, thereby effectively and properly disposing of the electronic data.

[0099] The foregoing description of embodiments is illustrative of various aspects/embodiments of the present invention. Various modifications to the invention can be made to the preferred embodiments by those skilled in the art without departing from the true spirit and scope of the invention as defined by the

[00100] appended claims. Accordingly, the scope of the present invention is defined by the appended claims rather than the foregoing description of embodiments.

What is claimed is: